

Latent Reward: LLM-Empowered Credit Assignment in Episodic Reinforcement Learning

Yun Qu*, Yuhang Jiang*, Boyuan Wang, Yixiu Mao, Cheems Wang, Chang Liu, Xiangyang Ji[†]

Tsinghua University

{qy22, jiangyh19, wangby22, myx21}@mails.tsinghua.edu.cn,
cheemswang@mail.tsinghua.edu.cn, liuchang2022@tsinghua.edu.cn, xyji@tsinghua.edu.cn

Abstract

Reinforcement learning (RL) often encounters delayed and sparse feedback in real-world applications, even with only episodic rewards. Previous approaches have made some progress in reward redistribution for credit assignment but still face challenges, including training difficulties due to redundancy and ambiguous attributions stemming from overlooking the multifaceted nature of mission performance evaluation. Hopefully, Large Language Model (LLM) encompasses fruitful decision-making knowledge and provides a plausible tool for reward redistribution. Even so, deploying LLM in this case is non-trivial due to the misalignment between linguistic knowledge and the symbolic form requirement, together with inherent randomness and hallucinations in inference. To tackle these issues, we introduce **LaRe**, a novel LLM-empowered symbolic-based decision-making framework, to improve credit assignment. Key to LaRe is the concept of the *Latent Reward*, which works as a multi-dimensional performance evaluation, enabling more interpretable goal attainment from various perspectives and facilitating more effective reward redistribution. We examine that semantically generated code from LLM can bridge linguistic knowledge and symbolic latent rewards, as it is executable for symbolic objects. Meanwhile, we design latent reward self-verification to increase the stability and reliability of LLM inference. Theoretically, reward-irrelevant redundancy elimination in the latent reward benefits RL performance from more accurate reward estimation. Extensive experimental results witness that LaRe (i) achieves superior temporal credit assignment to SOTA methods, (ii) excels in allocating contributions among multiple agents, and (iii) outperforms policies trained with ground truth rewards for certain tasks.

Introduction

Episodic reinforcement learning is dedicated to solving problems of receiving only episodic rewards, a frequently encountered situation in real-world applications of RL, such as autonomous driving (Kiran et al. 2021) and healthcare (Zeng et al. 2022). Credit assignment (Sutton et al. 2011; Zhang, Veeriah, and Whiteson 2020), which involves assessing the contributions of single-step decisions (Ren

et al. 2021), is challenging in episodic RL due to delayed and sparse feedback. Return decomposition (Arjona-Medina et al. 2019), which estimates proxy rewards by using state-action pairs to redistribute episodic rewards, has emerged in the literature as a promising direction to remedy this issue. Subsequent works often focus on model architectures (Liu et al. 2019; Widrich et al. 2021) or human-designed regression principles (Ren et al. 2021; Lin et al. 2024), overlooking the training difficulty posed by redundant information. Zhang et al. (2024b) attempted to address this redundancy by employing a causal approach to filter out reward-irrelevant features but still struggled with the lack of semantic interpretation.

A prominent observation in human problem-solving is that contribution assessments often encompass a range of qualitative and quantitative factors. For instance, soccer players’ performance is evaluated not only by goals scored but also by injury prevention and coordination. Similarly, the rewards designed in RL are commonly a combination of multiple factors (Todorov, Erez, and Tassa 2012; Qu et al. 2023). Previous methods (Arjona-Medina et al. 2019; Ren et al. 2021) mainly focus solely on the values of final returns without tapping into the multifaceted nature of performance evaluation, resulting in poor semantic interpretability and ambiguous credit assignment. Recently, the demonstrated capabilities of pre-trained LLM (Achiam et al. 2023) suggest that integrating its prior knowledge for improved credit assignment is a promising solution. However, the misalignment between LLM’s linguistic knowledge and the symbolic representations required for specific tasks poses significant challenges, while the inherent randomness and hallucinations in LLM inference further diminish its effectiveness (Peng et al. 2023; Carta et al. 2023).

Motivated by the urgent demand of depicting multifaceted performance evaluation, we propose a key concept for credit assignment, termed *Latent Reward*, where different dimensions capture various aspects of task performance while eliminating reward-irrelevant redundancy. We then devise a framework **LaRe**, which (i) derives semantically interpretable latent rewards by incorporating task-related priors from LLM and (ii) utilizes them to enhance reward decomposition. With the insight that semantically generated code can bridge linguistic knowledge in LLM and targets in symbolic form due to its executability for symbolic

*These authors contributed equally.

[†]Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

objects, LaRe presents a general paradigm for integrating LLM’s prior knowledge into symbolic tasks. Specifically, pre-trained LLM is instructed by standardized prompts to code encoding functions, which encode environment information into semantically interpretable latent rewards, eliminating the need for task-specific training. To alleviate the randomness and hallucinations in LLM reasoning, LaRe designs a self-verification mechanism for stable and reliable latent reward derivation.

Our **main contributions** are summarized as follows:

1. We propose the concept of Latent Reward with semantic interpretability and reveal the multifaceted nature of step-wise contributions by introducing it in the probabilistic model of episodic rewards, aligning with human preferences and reducing redundancy.
2. We devise a latent reward-based framework, LaRe, to leverage LLM’s task-related priors for more accurate and interpretable credit assignment, which paves a way for integrating LLM into symbolic-based decision-making.
3. We demonstrate the superiority of our method both theoretically and empirically and validate LLM’s effectiveness as a generalized information encoder for latent reward derivation in practical implementation.

The surprising phenomenon that LaRe outperforms policies trained with ground truth dense rewards for certain tasks highlights the significance of the semantically interpretable latent rewards derived through LLM’s reasoning capability. Our work reveals that merely fitting the final reward value, which primarily reflects overall performance, may be insufficient for effective reward decomposition. This suggests that RL can be further enhanced through multifaceted performance assessments informed by task-related priors.

Related Works

Reward Redistribution

Reward redistribution seeks to transform episodic rewards into immediate and dense proxy rewards \hat{r}_t , re-assigning credit for each state-action pair (Ren et al. 2021; Zhang et al. 2024b). Some previous methods focus on reward shaping (Ng, Harada, and Russell 1999; Hu et al. 2020) and intrinsic reward design (Pathak et al. 2017; Zheng et al. 2021). Return decomposition has emerged as a promising approach for tackling scenarios with severely delayed rewards. RUDDER (Arjona-Medina et al. 2019) analyzes the return-equivalent condition for invariant optimal policy and proposes return decomposition via a regression task. Subsequent works build on it by aligning demonstration sequences (Patil et al. 2020), using sequence modeling (Liu et al. 2019), or Hopfield networks (Widrich et al. 2021). Ren et al. (2021) propose randomized return decomposition to bridge between return decomposition (Efroni, Merlis, and Mannor 2021) and uniform reward redistribution (Gangwani, Zhou, and Peng 2020). Other redistribution principles have been adopted in recent works, such as causal treatment (Zhang et al. 2024b) and randomly cutting sub-trajectories (Lin et al. 2024). Recently, some methods have used attention-based approaches to decompose returns across time and agents in multi-agent settings (She,

Gupta, and Kochenderfer 2022; Xiao, Ramasubramanian, and Poovendran 2022; Chen et al. 2023). Despite significant progress, previous studies have neglected redundant reward-irrelevant features and the multifaceted nature of mission performance evaluation, which impede training and cause ambiguous attributions. While Zhang et al. (2024b) have acknowledged this issue to some extent, they focus solely on extracting reward-related state elements. In contrast, we propose the latent reward as a semantically interpretable multi-dimensional performance measurement and achieve reward-irrelevant redundancy elimination with task-related priors.

LLM-Empowered Decision Making

The remarkable capabilities of LLMs, as demonstrated across various downstream tasks (Touvron et al. 2023; Brown et al. 2020), underscores their potential as a promising solution for decision-making (Wang et al. 2023b). Some works focus on high-level control by employing LLMs as planners with predefined skills or APIs, which have proven highly successful (Liang et al. 2023; Yao et al. 2022; Shinn et al. 2023; Zhu et al. 2023; Wang et al. 2023a; Zhang et al. 2024a). However, when directly applied to low-level control without predefined skills, the misalignment between LLMs’ linguistic knowledge and the symbolic states and actions required for specific tasks poses a significant challenge (Peng et al. 2023; Qu et al. 2024b). Some works address this issue by constructing text-based environments but at the cost of considerable manual effort (Du et al. 2023; Carta et al. 2023). Recently, LLMs have been integrated with RL to enhance low-level control (Cao et al. 2024). Some approaches fine-tune LLMs as policies (Carta et al. 2023; Shi et al. 2024) or use LLM for history compression (Paischer et al. 2022). Other studies (Zhang et al. 2023; Su and Zhang 2023; Shukla et al. 2023) focus on goal-conditioned RL with LLMs as subgoal selectors, but these often require predefined skills or subgoals. We seek to leverage LLMs as tools to enhance RL, aligning with LLM-based reward design methods (Kwon et al. 2023; Song et al. 2023; Wang et al. 2024). However, our method ensures a more reliable and optimized use of LLM priors by strategically designing for improved response quality and integrating them into latent rewards during the training process for optimization rather than relying on unreliable direct use.

Preliminary

The environments in reinforcement learning are generally formulated by a Markov Decision Process (MDP; Bellman (1966)), which can be defined as a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \gamma, P, r \rangle$, where \mathcal{S} and \mathcal{A} denote the state space and action space with cardinalities $\|\mathcal{S}\|$ and $\|\mathcal{A}\|$, respectively. $\gamma \in [0, 1)$ is the discount factor. $P(s'|s, a)$ represents the environment’s state transition distribution, and $r(s, a)$ denotes the reward function. The goal of reinforcement learning is to find an optimal policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes the expected cumulative rewards with the initial state distribution η and episode length T , which is expressed as $J(\pi) = \mathbb{E} \left[\sum_{t=1}^T \gamma^t r(s_t, \pi(s_t)) \mid s_0 \sim \eta, s_{t+1} \sim P(\cdot | s_t, \pi(s_t)) \right]$.

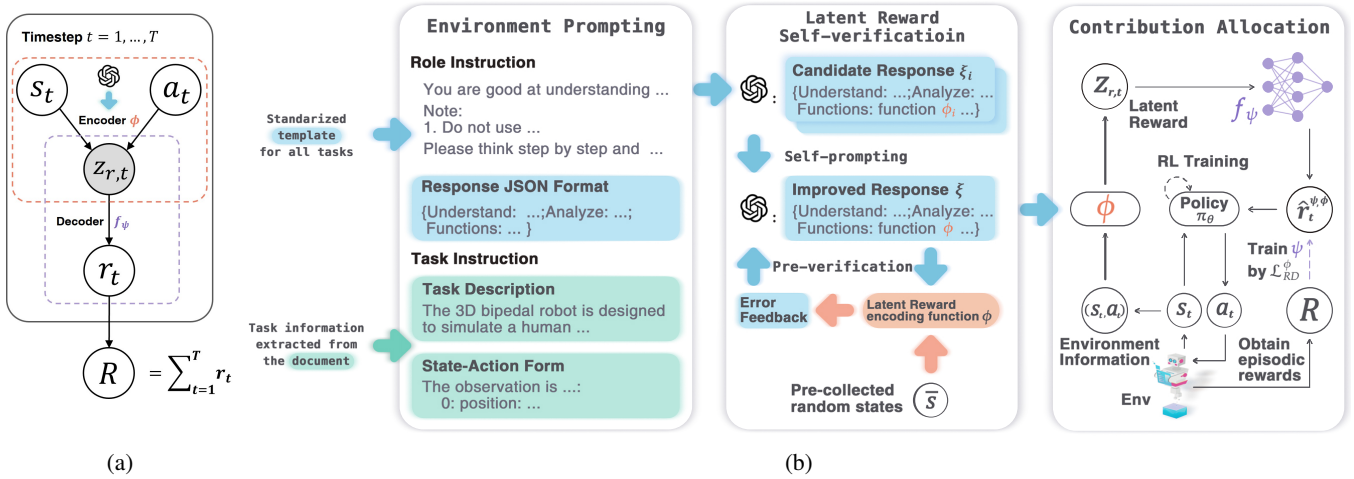


Figure 1: Overview of LaRe. (a) The probabilistic model of the episodic reward with the latent reward $z_{r,t}$ introduced as the implicit variable. (b) The LaRe framework consists of three main components: (1) Environment Prompting: the task information is incorporated into a standardized prompt for LLM instructions (details are in Appendix A). (2) Latent Reward Self-verification: during the self-prompting phase, LLM generates n candidate responses $\{\xi_i\}_{i=1}^n$ and synthesizes an improved response ξ . In the pre-verification phase, the executability of the function ϕ is verified with pre-collected random states \bar{s} ; (3) Contribution Allocation: latent rewards $z_{r,t}$ are derived by ϕ and used to estimate proxy rewards via the reward decoder model f_ψ .

Real-world scenarios often pose challenges such as delayed and sparse feedback (Ke et al. 2018; Han et al. 2022). An extreme case is episodic RL, where only a non-zero reward $R(\tau)$ at the end of each trajectory τ is received (Ren et al. 2021). The goal of episodic reinforcement learning is to maximize the expected episodic rewards, i.e., $J_{ep}(\pi) = \mathbb{E}[R(\tau)|s_0 \sim \eta, a_t \sim \pi(\cdot|s_t), \tau = \langle s_0, a_0, s_1, \dots, s_T \rangle]$. A common assumption in episodic RL is the existence of a sum-form decomposition of the episodic rewards, i.e., $R(\tau) = \sum_{t=1}^T r(s_t, a_t)$ (Zhang et al. 2024b).

Latent Reward

This section elaborates on LaRe’s motivation and implementation. We explain the rationale behind the *Latent Reward* and analyze the underlying probabilistic model. We propose a framework *LaRe* that leverages LLM’s reasoning and generalization capabilities while addressing the challenges of its application to incorporate task-related prior for reliably deriving the latent reward. We theoretically prove that by reducing reward-irrelevant redundancy, the latent reward enhances reward modeling and improves RL performance.

Motivation

In human endeavors, individual contributions are typically assessed from multiple angles for a comprehensive evaluation. However, current research on episodic credit assignment often focuses solely on regressing the final reward values (Arjona-Medina et al. 2019; Efroni, Merlis, and Mannor 2021), overlooking that rewards are derived from the evaluation of various implicit factors, such as costs and efficiency. Inspired by the intrinsic need to evaluate task performance from multiple perspectives, we propose the concept of the

Latent Reward. Conceptually, the different dimensions of latent reward capture various aspects of task performance.

Formally, the reward r is a projection of the latent reward z_r from a space \mathcal{D} with cardinality $\|\mathcal{D}\|$ onto the real number field \mathbb{R} . A function $f : \mathcal{D} \rightarrow \mathbb{R}$ should exist such that each reward in the reward codomain has at least one latent reward encoding. With the introduction of the latent reward, as illustrated in Figure 1a, we construct a new probabilistic model of the episodic reward, revealing the multifaceted nature of the step-wise contribution, which better serves RL training. We have,

$$\begin{aligned}
 p(R|s_{1:T}, a_{1:T}) &= \int p(R, r_{1:T}, z_{r,1:T}|s_{1:T}, a_{1:T}) dz dr \\
 &= \int \left[\prod_{t=1}^T \underbrace{p(r_t|z_{r,t})}_{\text{decoder } f} \underbrace{p(z_{r,t}|s_t, a_t)}_{\text{encoder } \phi} \right] p(R|r_{1:T}) dz dr \quad (1)
 \end{aligned}$$

where the $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{D}$ is the function deriving the latent reward from environment information. Intuitively, the latent reward’s multiple dimensions are obtained by compressing environmental information based on prior knowledge, thus acting as an information bottleneck (Tishby, Pereira, and Bialek 2000) tailored to the task objectives.

Compared to directly estimating step-wise rewards from raw states, the latent reward offers significant advantages in interpretability, as each dimension reflects a specific aspect of task performance. Additionally, in episodic RL, where only the return of an episode provides weak signals, directly modeling rewards can be challenging. Learning from latent rewards better aligns with task objectives and simplifies network training by reducing reward-irrelevant redundancy.

A naive approach is to obtain the latent reward via an information bottleneck method, which suffers from limited

Algorithm 1: LaRe

Input: LLM \mathcal{M} , task information $task$, role instruction $role$, candidate responses number n , pre-collected random state-action pairs \bar{s} , max episodes \mathcal{N}^{max}

Output: policy network π_θ , reward decoder model f_ψ

- 1: Initialize the policy network parameter θ , the reward decoder model parameter ψ , and the replay buffer \mathcal{B} .
- 2: Obtain response ξ by executing Eq. (3) and Eq. (4).
- 3: Repeat Eq. (5) until obtaining an executable ϕ .
- 4: **for** $episode = 1$ to \mathcal{N}^{max} **do**
- 5: Sample a trajectory τ using current policy.
- 6: $\mathcal{B} \leftarrow \mathcal{B} \cup \{\tau\}$. Sample a batch $B = \{\tau_i\}_{i=1}^{|B|}$ from \mathcal{B} .
- 7: Estimate latent reward enhanced return decomposition loss $\mathcal{L}_{RD}^\phi(\psi)$ with Eq. (6) and update reward decoder model f_ψ

$$\psi \leftarrow \psi - \alpha \nabla_\psi \mathcal{L}_{RD}^\phi(\psi) \quad (2)$$

- 8: Perform policy optimization using any RL algorithm with predicted proxy rewards $\hat{r}^{\psi, \phi} = f_\psi(\phi(s, a))$.
 - 9: **end for**
-

linguistic interpretability and high computational costs due to separate encoder training for each task. In contrast, LLM’s pre-training has captured more compact representations in the form of tokens, facilitating better cross-task generalization. Therefore, leveraging LLM’s prior knowledge enables more efficient extraction of interpretable and multifaceted task performance metrics, the latent reward, from the redundant environmental information.

Framework

Leveraging LLM’s prior knowledge and reasoning capabilities to derive latent rewards for credit assignment presents three main challenges: (1) instructing LLM to derive latent rewards for various tasks with minimal information and effort, (2) addressing the linguistic-symbolic misalignment while mitigating randomness and hallucinations in LLM inference to derive symbolic latent rewards reliably, and (3) applying latent rewards to enhance contribution allocation at each timestep. This section introduces three specifically designed components in the proposed LaRe, as demonstrated in Fig. 1b and Algorithm 1:

Environment Prompting. To instruct LLM, we design standardized prompts easily transferable across environments, which consist of a templated role instruction ($role$) and specific task instruction ($task$), as shown in Fig. 1b. The role instruction is consistent across tasks and guides LLM to think in a predefined manner: understand the task and state \rightarrow identify reward-related factors \rightarrow generate the latent reward encoding function. Only the necessary task description and state forms for a specific task are required, which can be easily extracted from the task document. The task description mainly includes the environment profile and task objective. The state forms detail the meanings of dimensions in the state space. Our design significantly reduces the burden of labor-intensive prompt engineering across tasks.

Latent Reward Self-verification. Since LLM’s knowledge is encoded in language while underlying tasks are represented by symbolic states, this misalignment impedes LLM’s direct application. To effectively integrate LLM, we propose generating the latent reward encoding function using LLM’s coding capabilities. The rationale is that semantically generated code can bridge the gap between linguistic knowledge and symbolic latent rewards, as its execution is symbolic and tailored to specific tasks, as previously confirmed (Wang et al. 2024). Given the inherent randomness and hallucinations in LLM inference, inspired by recent work (Shinn et al. 2023; Ma et al. 2023), we propose a latent reward LLM generation process with self-verification, which includes *self-prompting* and *pre-verification* to enhance stability and reliability.

In the self-prompting phase, LLM \mathcal{M} firstly generates n candidate responses, each including a code implementation of the latent reward encoding function:

$$\xi_1, \xi_2, \dots, \xi_n \leftarrow \mathcal{M}(task, role) \quad (3)$$

These candidate responses are then fed into the prompt, and LLM is prompted to summarize an improved response:

$$\xi \leftarrow \mathcal{M}(task, role, \xi_{1..n}) \quad (4)$$

Regarding pre-verification, leveraging the standardized response template, the latent reward encoding function ϕ can be easily extracted from the response ξ , which takes in a state-action pair s, a and outputs a latent reward $z_r = \phi(s) = [z_r^1, \dots, z_r^d]$. We then verify ϕ with pre-collected random state-action pairs \bar{s} and provide error feedback to LLM until ϕ is executable:

$$err \leftarrow verify(\phi, \bar{s}); \xi \leftarrow \mathcal{M}(task, role, \xi_{1..n}, err) \quad (5)$$

Self-verification significantly improves response quality by reducing randomness in identifying latent rewards and ensuring code executability. LLM’s clear linguistic responses and transparent thought processes provide high interpretability, facilitating human evaluation and manual intervention. Empirical results demonstrate that our framework achieves satisfactory results without requiring multi-iteration evolutionary optimization (Ma et al. 2023).

Contribution Allocation. Building on the latent reward encoding function, we adopt a latent reward enhanced return decomposition, implemented based on Efroni, Merlis, and Mannor (2021). Let f_ψ be a neural network decoder parameterized by ψ . The new objective of reward modeling can be formulated as:

$$\min_\psi \mathcal{L}_{RD}^\phi(\psi) = \mathbb{E}_{\tau \sim D} \left[\left(R(\tau) - \sum_{t=1}^T f_\psi(\phi(s_t, a_t)) \right)^2 \right] \quad (6)$$

Proxy rewards, $\hat{r}^{\psi, \phi} = f_\psi(\phi(s, a))$, derived from latent rewards, are incorporated into the RL training process. Leveraging the enhanced temporal credit assignment enabled by the latent reward’s multifaceted nature, these rewards improve RL training performance by alleviating the issue of delayed and sparse feedback.

Additionally, we empirically find that the latent reward enhances credit assignment among agents. This well




Figure 2: Average episode return for tasks with different state space dimensions in MuJoCo. Notably, **TD3-DR** is trained with dense rewards.

matches the intuition, as evaluating agents within a team is also a form of multifaceted credit assignment. Consequently, our method provides a practical solution for episodic multi-agent RL, with reduced computational costs and improved performance, making it well-suited for real-world scenarios.

In implementations, we use GPT-4o from OpenAI API, with prompt details provided in Appendix A (Qu et al. 2024a). In practice, we have set the random variables deterministically for the sake of convenience, which is a common setting in previous works (Arjona-Medina et al. 2019).

Analysis

LLM-empowered latent rewards retain semantic interpretability while reducing reward-irrelevant redundancy, which is theoretically proven to boost RL performance by learning a better reward model than the state-based methods.

Previous works commonly minimize the least squares error between the episodic rewards and the sum of predicted proxy rewards $\hat{r}(s_t, a_t)$ to learn reward models with raw states as inputs (Ren et al. 2021). The surjective function $\phi(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{D}, \|\mathcal{D}\| < \|\mathcal{S}\| \|\mathcal{A}\|$ reduces redundant, reward-irrelevant features from the state-action space. Theoretically, built upon Efroni, Merlis, and Mannor (2021), assuming access to a latent reward function ϕ that satisfies $\exists f^*, s.t., r = \hat{r} = f^*(\phi(s, a))$, we derive a more precise concentration bound for estimating r and a tighter RL regret bound compared to the case without the latent reward. Please refer to Appendix B for the proof.

Proposition 1 (Tighter Concentration Bound of Reward).

Let $\lambda > 0$ and $A_k^\phi \stackrel{\text{def}}{=} (H_k^\phi)^T H_k^\phi + \lambda I_{\|\mathcal{D}\|}$. For any $\delta \in (0, 1)$, with probability greater than $1 - \delta/10$ uniformly for all episode indexes $k \geq 0$, it holds that

$$\|r - \hat{r}_k^\phi\|_{A_k^\phi} \leq \sqrt{\frac{1}{4} T \|\mathcal{D}\| \log\left(\frac{1 + kT^2/\lambda}{\delta/10}\right)} + \sqrt{\lambda \|\mathcal{D}\|} \stackrel{\text{def}}{=} l_k^\phi < l_k.$$

Proposition 2 (Tighter Regret Bound).

For any $\delta \in (0, 1)$ and all episode numbers $K \in \mathbb{N}^+$, the regret of RL $\rho^\phi(K) \stackrel{\text{def}}{=} \sum_{k=1}^K (V^* - V^{\phi, \pi_k})$ holds with probability greater than $1 - \delta$ that,

$$\rho^\phi(K) \leq \mathcal{O}\left(T \|\mathcal{D}\| \sqrt{K} \log\left(\frac{KT}{\delta}\right)\right) < \mathcal{O}\left(T \|\mathcal{S}\| \|\mathcal{A}\| \sqrt{K} \log\left(\frac{KT}{\delta}\right)\right).$$

The concentration bound reflects the performance of the reward model by quantifying the distance between proxy rewards \hat{r}_k^ϕ and true rewards r , while the regret quantifies RL performance. Proposition 1 and 2 show that these bounds are proportional to $\|\mathcal{D}\|$, which are lower than the bound with raw state-action space. Overall, the latent reward improves reward function learning and boosts RL performance.

Experiments

We evaluate LaRe¹ on two widely used benchmarks in both single-agent and multi-agent settings: MuJoCo locomotion benchmark (Todorov, Erez, and Tassa 2012) and Multi-Agent Particle Environment (MPE) (Lowe et al. 2017). Additionally, we perform ablation studies and further analyses to validate LaRe’s components and assess its properties.

Experimental Setups

For MuJoCo, we adopt four tasks from Gymnasium (Towers et al. 2023). For MPE, we employ six tasks from two scenarios, *Cooperative-Navigation (CN)* and *Predator-Prey (PP)*, featuring varying numbers of agents (6, 15, 30), which are based on Chen et al. (2023) with minor modifications to provide individual rewards to each agent at every step. All tasks are episodic, with a single non-zero episodic reward, equivalent to the cumulative rewards. Thus, multi-agent tasks require both temporal and inter-agent credit assignment. Moreover, we evaluate LaRe in more complex scenarios from SMAC (Samvelyan et al. 2019) and a **newly designed task, Triangle Area**, in Appendix D and E.

We compare LaRe with SOTA return decomposition baseline algorithms: RD (Efroni, Merlis, and Mannor 2021), IRCR (Gangwani, Zhou, and Peng 2020), Diaster (Lin et al. 2024), RRD and RRD.unbiased (Ren et al. 2021), as well as those designed for multi-agent settings: AREL (Xiao, Ramasubramanian, and Poovendran 2022) and STAS (Chen et al. 2023). The introduction and implementation details of these baselines are provided in Appendix C.

LaRe is compatible with various RL algorithms, and we adopt TD3 (Fujimoto, Hoof, and Meger 2018) for single-agent and IPPO (Yu et al. 2022) for multi-agent as the base algorithm, consistent with prior works (Ren et al. 2021;

¹Our code is available at <https://github.com/thu-rlab/LaRe>

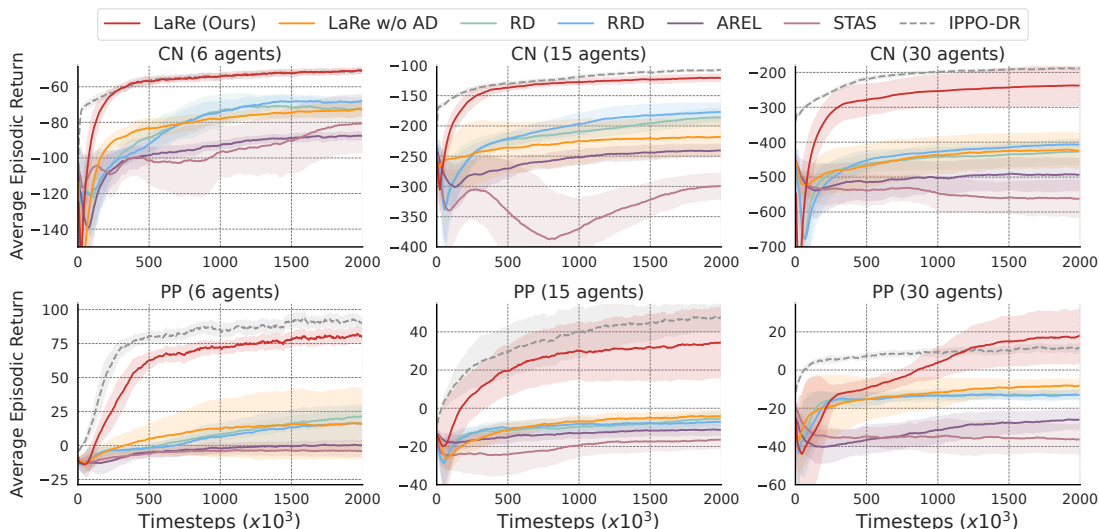


Figure 3: Average episode return for tasks with a varied number of agents in MPE. Notably, **IPPO-DR** is trained with dense rewards and **LaRe w/o AD** represents LaRe without credit assignment among agents.

Chen et al. 2023). Each algorithm runs on five random seeds, with the mean performance and standard deviation reported. Further details and results are available in the Appendix.

The Superiority of LaRe

Single-Agent. To verify the compatibility of our method with various return decomposition algorithms, we implement two variants, **LaRe-RD** and **LaRe-RRDu**, based on RD and RRD-unbiased, respectively. As shown in Fig. 2, the poor performance of TD3 and IRCR highlights the importance of assigning individual credits. Our method, LaRe, consistently outperforms SOTA baselines on MuJoCo tasks, demonstrating higher sample efficiency and better convergence. Both variants of LaRe surpass the corresponding baselines, highlighting the efficacy of semantically interpretable latent rewards in credit assignment. The effectiveness of LaRe in tasks with large state spaces significantly supports our analysis, underscoring the significance of redundancy elimination with task-related priors in the latent reward and explaining the poor performance of baselines.

Multi-Agent. Fig. 3 depicts comprehensive comparisons between LaRe and various baselines in MPE. LaRe is implemented based on RD and demonstrates superior performance across tasks with different numbers of agents compared to all SOTA baselines, confirming the efficacy of latent rewards in temporal credit assignment on multi-agent tasks. We also include a variant **LaRe w/o AD** (without agent decomposition), where the proxy rewards of different agents are averaged at the same time step. The significant performance drop highlights the necessity of credit assignment at the agent level and the effectiveness of LaRe in this regard. We believe the semantically interpretable latent rewards account for this since assessing different agents’ contributions is also intuitively a form of multifaceted credit assignment. AREL and STAS perform relatively poorly, particularly as the number of agents increases, likely because

reward-irrelevant items in the original state interfere significantly with attention-based credit assignment.

Comparable with Dense Rewards. We include **TD3-DR** and **IPPO-DR** in MuJoCo and MPE, respectively, training with ground truth dense rewards. Remarkably, LaRe’s performance is comparable to or even exceeds theirs despite not relying on manually designed rewards. The reason is that while ground true rewards reflect agents’ performance levels, overall stability is still affected by implicit factors like costs and efficiency, which are adequately captured by our proposed LLM-based latent reward. This finding emphasizes leveraging task-related prior information for multifaceted performance evaluation can further enhance RL performance beyond merely relying on final reward values.

Delving into Latent Rewards

We conduct experiments to analyze the specific nature of the latent rewards and the reason for their superior performance.

Semantic Analysis of Multifaceted Measurement. We analyze the LLM-generated latent reward functions and use *HumanoidStandup-v4* as an instance. The task objective is to have the humanoid robot stand up and maintain balance by applying torques to hinges (Towers et al. 2023). As shown in Fig. 4(b), LLM demonstrates a correct understanding of the task and derives latent rewards as interpretable performance measures across multiple dimensions, such as height and safe control, which align with the ground truth (GT) reward function. Additionally, LLM considers stability, which better aligns with the task’s objectives, further elucidating its superior performance compared to baselines with dense rewards. Further details can be found in Appendix A.

Reduced Reward-irrelevant Redundancy. We calculate the Pearson correlation coefficient (Cohen et al. 2009) between each dimension of original states or LLM-generated latent rewards and ground truth dense rewards. As shown in Table 1, latent rewards are tighter correlated with ground




Figure 4: (a) The task *HumanoidStandup-v4* aims to make the humanoid stand up and maintain balance. (b) LLM-generated latent rewards additionally consider implicit factors affecting stability compared to the ground truth rewards. (c) Comparison between LaRe and RD on the competitive *Predator-Prey (6 agents)* task. ‘X vs Y’ represents the condition where X controls preys and Y controls predators. LaRe outperforms RD when directly pitted against it.

Tasks	corr (dims)		exe_rate	
	States	Latent Rewards	w/o PV	w/ PV
<i>CN (6 agents)</i>	0.02 (26)	0.50 (5.6)	0%	100%
<i>PP (6 agents)</i>	0.01 (28)	0.12 (5.4)	20%	100%
<i>HalfCheetah-v4</i>	0.22 (17)	0.53 (4.8)	40%	100%
<i>HumanoidStandup-v4</i>	0.20 (376)	0.49 (5.6)	40%	100%

Table 1: `corr` denotes the average Pearson correlation coefficient. `dims` represents the average number of dimensions of original states or latent rewards. Additionally, we record the average execution rate `exe_rate` of LLM-generated latent reward functions without pre-verification (w/o PV).

truth rewards across tasks. Meanwhile, latent rewards’ dimensions are significantly fewer than those of original states. The results confirm that the latent reward reduces the reward-irrelevant redundancy with task-related priors, improving reward prediction, as shown in Appendix F.5.

Algorithm Agnostic. Notably, latent rewards for estimating proxy rewards are transferable to various RL backbones. This property ensures LaRe’s application prospects, opening up possibilities to combine with real-world approaches. We conduct detailed experiments in Appendix F.3.

Compatible with Heterogeneous Agents. Latent rewards can help re-assign credits among heterogeneous agents, even in competitive scenarios. Like Lowe et al. (2017), we jointly train policies for competitive predators and preys in task *Predator-Prey*. We have the policies trained by LaRe and RD respectively compete, with preys and predators controlled by different ones. As shown in Fig. 4(c), LaRe learns superior policies for both predators and preys compared to RD, suggesting enhanced credit assignment in competitive multi-agent scenarios. This advantage can be attributed to the multifaceted nature of the latent reward.

Ablation Studies

Reward Attributes in Latent Reward. To distinguish latent rewards from mere state representation, we conduct an ablation study by removing the reward decoder model, termed “LaRe w/o RM”, which estimates proxy re-

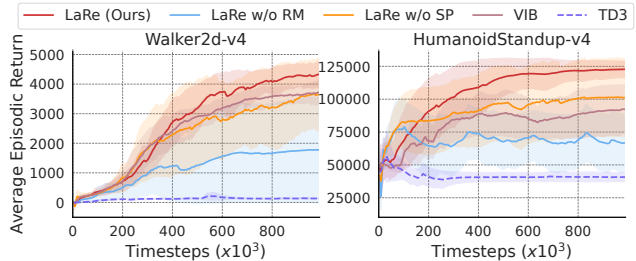


Figure 5: Ablation studies of the reward model and the proposed self-verified LLM generation, as well as comparisons of LaRe with the variational information bottleneck.

wards by summarizing latent rewards with a sign: $\hat{r}_{sign} = \sum_{i=1}^d \text{sign}(z_r^i) \cdot z_r^i$. The signs are obtained by minimizing the estimation loss between episodic rewards and the sum of proxy rewards. As shown in Figure 5, this significant simplification outperforms the baseline with episodic rewards (TD3), confirming that the latent reward possesses genuine reward attributes rather than just representing states.

Self-Verification. We propose Self-prompting (SP) and pre-verification (PV) to reduce randomness and hallucinations in LLM inference. The superior performance shown in Fig. 5 indicates that SP effectively reduces randomness in LLM inference, resulting in improved LLM responses and RL performance. In Table. 1, we compare the execution rate of LLM-generated latent reward encoding functions with and without pre-verification, highlighting the significance of pre-verification. The results validate the efficacy of our designs for integrating LLMs into RL tasks. Appendix F.1 examines the impact of the number of candidate responses.

Variational Information Bottleneck. We propose an alternative method to utilize the Variational Information Bottleneck (VIB) (Alemi et al. 2017) to derive latent rewards. However, as shown in Fig. 5 and Appendix F.2, LaRe significantly outperforms VIB, which underscores the importance of LLMs serving as generalized encoders of environmental information to integrate task-related priors effectively. We compare LaRe with LLM reward design in Appendix F.4.

Conclusion

In this work, we present LaRe, a LLM-empowered framework for credit assignment in episodic reinforcement learning with task-related prior. The framework is centered on the latent reward, whose dimensions reflect distinct aspects of task performance evaluation. We utilize LLM’s coding abilities to address the linguistic-symbolic misalignment in integrating LLM into RL tasks and propose self-verification to ensure reliable LLM inference. This work (i) addresses previously overlooked research questions, including training difficulties caused by redundancy and the multifaceted nature of mission performance evaluation; (ii) develops a practical algorithm that achieves superior performance both theoretically and empirically; (iii) advances the integration of LLM prior knowledge into RL through semantically interpretable multifaceted performance evaluation.

Acknowledgments

This work was supported by the National Key R&D Program of China under Grant 2018AAA0102801.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2017. Deep Variational Information Bottleneck. In *ICLR*.
- Arjona-Medina, J. A.; Gillhofer, M.; Widrich, M.; Unterthiner, T.; Brandstetter, J.; and Hochreiter, S. 2019. Rudder: Return decomposition for delayed rewards. *NeurIPS*.
- Bellman, R. 1966. Dynamic programming. *science*, 153(3731): 34–37.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *NeurIPS*, 33: 1877–1901.
- Cao, Y.; Zhao, H.; Cheng, Y.; Shu, T.; Liu, G.; Liang, G.; Zhao, J.; and Li, Y. 2024. Survey on large language model-enhanced reinforcement learning: Concept, taxonomy, and methods. *arXiv preprint arXiv:2404.00282*.
- Carta, T.; Romac, C.; Wolf, T.; Lamprier, S.; Sigaud, O.; and Oudeyer, P.-Y. 2023. Grounding large language models in interactive environments with online reinforcement learning. *arXiv preprint arXiv:2302.02662*.
- Chen, S.; Zhang, Z.; Yang, Y.; and Du, Y. 2023. STAS: Spatial-Temporal Return Decomposition for Multi-agent Reinforcement Learning. *arXiv preprint arXiv:2304.07520*.
- Cohen, I.; Huang, Y.; Chen, J.; Benesty, J.; Benesty, J.; Chen, J.; Huang, Y.; and Cohen, I. 2009. Pearson correlation coefficient. *Noise reduction in speech processing*, 1–4.
- Du, Y.; Watkins, O.; Wang, Z.; Colas, C.; Darrell, T.; Abbeel, P.; Gupta, A.; and Andreas, J. 2023. Guiding pre-training in reinforcement learning with large language models. *arXiv preprint arXiv:2302.06692*.
- Efroni, Y.; Merlis, N.; and Mannor, S. 2021. Reinforcement learning with trajectory feedback. In *AAAI*.
- Fujimoto, S.; Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. *ICML*.
- Gangwani, T.; Zhou, Y.; and Peng, J. 2020. Learning guidance rewards with trajectory-space smoothing. *NeurIPS*.
- Han, B.; Ren, Z.; Wu, Z.; Zhou, Y.; and Peng, J. 2022. Off-policy reinforcement learning with delayed rewards. In *ICML*, 8280–8303. PMLR.
- Hu, Y.; Wang, W.; Jia, H.; Wang, Y.; Chen, Y.; Hao, J.; Wu, F.; and Fan, C. 2020. Learning to utilize shaping rewards: A new approach of reward shaping. *NeurIPS*.
- Ke, N. R.; ALIAS PARTH GOYAL, A. G.; Bilaniuk, O.; Binas, J.; Mozer, M. C.; Pal, C.; and Bengio, Y. 2018. Sparse attentive backtracking: Temporal credit assignment through reminding. *NeurIPS*, 31.
- Kiran, B. R.; Sobh, I.; Talpaert, V.; Mannion, P.; Al Sallab, A. A.; Yogamani, S.; and Pérez, P. 2021. Deep reinforcement learning for autonomous driving: A survey. *IEEE T-ITS*.
- Kwon, M.; Xie, S. M.; Bullard, K.; and Sadigh, D. 2023. Reward Design with Language Models. In *ICLR*.
- Liang, J.; Huang, W.; Xia, F.; Xu, P.; Hausman, K.; Ichter, B.; Florence, P.; and Zeng, A. 2023. Code as policies: Language model programs for embodied control. In *ICRA*.
- Lin, H.; Wu, H.; Zhang, J.; Sun, Y.; Ye, J.; and Yu, Y. 2024. Episodic Return Decomposition by Difference of Implicitly Assigned Sub-trajectory Reward. In *AAAI*.
- Liu, Y.; Luo, Y.; Zhong, Y.; Chen, X.; Liu, Q.; and Peng, J. 2019. Sequence modeling of temporal credit assignment for episodic reinforcement learning. *arXiv preprint arXiv:1905.13420*.
- Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Abbeel, O. P.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *NeurIPS*.
- Ma, Y. J.; Liang, W.; Wang, G.; Huang, D.-A.; Bastani, O.; Jayaraman, D.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2023. Eureka: Human-level reward design via coding large language models. *arXiv preprint arXiv:2310.12931*.
- Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, 278–287.
- Paischer, F.; Adler, T.; Patil, V.; Bitto-Nemling, A.; Holzleitner, M.; Lehner, S.; Eghbal-Zadeh, H.; and Hochreiter, S. 2022. History compression via language models in reinforcement learning. In *ICML*, 17156–17185. PMLR.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *ICML*.
- Patil, V. P.; Hofmarcher, M.; Dinu, M.-C.; Dorfer, M.; Blies, P. M.; Brandstetter, J.; Arjona-Medina, J. A.; and Hochreiter, S. 2020. Align-rudder: Learning from few demonstrations by reward redistribution. *arXiv preprint arXiv:2009.14108*.
- Peng, S.; Hu, X.; Yi, Q.; Zhang, R.; Guo, J.; Huang, D.; Tian, Z.; Chen, R.; Du, Z.; Guo, Q.; et al. 2023. Self-driven Grounding: Large Language Model Agents with Automatic Language-aligned Skill Learning. *arXiv preprint arXiv:2309.01352*.
- Qu, Y.; Jiang, Y.; Wang, B.; Mao, Y.; Wang, C.; Liu, C.; and Ji, X. 2024a. Latent Reward: LLM-Empowered Credit Assignment in Episodic Reinforcement Learning. *arXiv preprint arXiv:2412.11120*.
- Qu, Y.; Wang, B.; Jiang, Y.; Shao, J.; Mao, Y.; Wang, C.; Liu, C.; and Ji, X. 2024b. Choices are more important than efforts: Llm enables efficient multi-agent exploration. *arXiv preprint arXiv:2410.02511*.
- Qu, Y.; Wang, B.; Shao, J.; Jiang, Y.; Chen, C.; Ye, Z.; Liu, L.; Feng, Y. J.; Lai, L.; Qin, H.; et al. 2023. Hokoff: Real Game Dataset from Honor of Kings and its Offline Reinforcement Learning Benchmarks. In *NeurIPS*.
- Ren, Z.; Guo, R.; Zhou, Y.; and Peng, J. 2021. Learning long-term reward redistribution via randomized return decomposition. *arXiv preprint arXiv:2111.13485*.

- Samvelyan, M.; Rashid, T.; de Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G. J.; Hung, C.-M.; Torr, P. H. S.; Foerster, J.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. *CoRR*, abs/1902.04043.
- She, J.; Gupta, J. K.; and Kochenderfer, M. J. 2022. Agent-time attention for sparse rewards multi-agent reinforcement learning. *arXiv preprint arXiv:2210.17540*.
- Shi, R.; Liu, Y.; Ze, Y.; Du, S. S.; and Xu, H. 2024. Unleashing the Power of Pre-trained Language Models for Offline Reinforcement Learning. In *ICLR*.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K. R.; and Yao, S. 2023. Reflexion: Language agents with verbal reinforcement learning. In *NeurIPS*.
- Shukla, Y.; Gao, W.; Sarathy, V.; Velasquez, A.; Wright, R.; and Sinapov, J. 2023. LgTS: Dynamic Task Sampling using LLM-generated sub-goals for Reinforcement Learning Agents. *arXiv preprint arXiv:2310.09454*.
- Song, J.; Zhou, Z.; Liu, J.; Fang, C.; Shu, Z.; and Ma, L. 2023. Self-Refined Large Language Model as Automated Reward Function Designer for Deep Reinforcement Learning in Robotics. *arXiv preprint arXiv:2309.06687*.
- Su, J.; and Zhang, Q. 2023. Subgoal Proposition Using a Vision-Language Model. In *CoRL Workshop on LEAP*.
- Sutton, R. S.; Modayil, J.; Delp, M.; Degris, T.; Pilarski, P. M.; White, A.; and Precup, D. 2011. Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *AAMAS*.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *IROS*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Towers, M.; Terry, J. K.; Kwiatkowski, A.; Balis, J. U.; Cola, G. d.; Deleu, T.; Goulão, M.; Kallinteris, A.; KG, A.; Krimmel, M.; Perez-Vicente, R.; Pierré, A.; Schulhoff, S.; Tai, J. J.; Shen, A. T. J.; and Younis, O. G. 2023. Gymnasium.
- Wang, B.; Qu, Y.; Jiang, Y.; Shao, J.; Liu, C.; Yang, W.; and Ji, X. 2024. LLM-Empowered State Representation for Reinforcement Learning. *arXiv preprint arXiv:2407.13237*.
- Wang, G.; Xie, Y.; Jiang, Y.; Mandlekar, A.; Xiao, C.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2023a. Voyager: An open-ended embodied agent with large language models. *TMLR*.
- Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; et al. 2023b. A survey on large language model based autonomous agents. *arXiv preprint arXiv:2308.11432*.
- Widrich, M.; Hofmarcher, M.; Patil, V. P.; Bitto-Nemling, A.; and Hochreiter, S. 2021. Modern hopfield networks for return decomposition for delayed rewards. In *Deep RL Workshop NeurIPS 2021*.
- Xiao, B.; Ramasubramanian, B.; and Poovendran, R. 2022. Agent-temporal attention for reward redistribution in episodic multi-agent reinforcement learning. *arXiv preprint arXiv:2201.04612*.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.
- Yu, C.; Velu, A.; Vinitzky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *NeurIPS*, 35: 24611–24624.
- Zeng, J.; Shao, J.; Lin, S.; Zhang, H.; Su, X.; Lian, X.; Zhao, Y.; Ji, X.; and Zheng, Z. 2022. Optimizing the dynamic treatment regime of in-hospital warfarin anticoagulation in patients after surgical valve replacement using reinforcement learning. *JAMIA*.
- Zhang, D.; Chen, L.; Zhang, S.; Xu, H.; Zhao, Z.; and Yu, K. 2024a. Large language models are semi-parametric reinforcement learning agents. *NeurIPS*, 36.
- Zhang, J.; Zhang, J.; Pertsch, K.; Liu, Z.; Ren, X.; Chang, M.; Sun, S.-H.; and Lim, J. J. 2023. Bootstrap your own skills: Learning to solve new tasks with large language model guidance. *arXiv preprint arXiv:2310.10021*.
- Zhang, S.; Veeriah, V.; and Whiteson, S. 2020. Learning retrospective knowledge with reverse reinforcement learning. *NeurIPS*, 33: 19976–19987.
- Zhang, Y.; Du, Y.; Huang, B.; Wang, Z.; Wang, J.; Fang, M.; and Pechenizkiy, M. 2024b. Interpretable reward redistribution in reinforcement learning: a causal approach. *NeurIPS*, 36.
- Zheng, L.; Chen, J.; Wang, J.; He, J.; Hu, Y.; Chen, Y.; Fan, C.; Gao, Y.; and Zhang, C. 2021. Episodic multi-agent reinforcement learning with curiosity-driven exploration. *NeurIPS*, 34: 3757–3769.
- Zhu, X.; Chen, Y.; Tian, H.; Tao, C.; Su, W.; Yang, C.; Huang, G.; Li, B.; Lu, L.; Wang, X.; et al. 2023. Ghost in the Minecraft: Generally Capable Agents for Open-World Environments via Large Language Models with Text-based Knowledge and Memory. *arXiv preprint arXiv:2305.17144*.